



## **Avere OS 4.6.2.5.C11 Release Notes**

2018-05-21

This page is intentionally left blank.

## Copyright Information

Copyright © 2009-2018 Avere Systems, Inc. All rights reserved. Specifications subject to change without notice.

No part of this document covered by copyright may be reproduced in any form or by any means – graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system – without prior written permission of the copyright owner.

The product described in this document may be protected by one or more U.S. patents, foreign patents, or pending applications.

## Table of Contents

[New in Avere OS 4.6.2.5.C10](#)

[New in Avere OS 4.6.2.5.C9](#)

[New in Avere OS 4.6.2.5.C8](#)

[New in Avere OS 4.6.2.5.C7](#)

[New in Avere OS 4.6.2.5.C6](#)

[New in Avere OS 4.6.2.5.C5](#)

[New in Avere OS 4.6.2.5.C4](#)

[New in Avere OS 4.6.2.5.C3](#)

[New in Avere OS 4.6.2.5.C2](#)

[New in Avere OS 4.6.2.5.C1](#)

[New in Avere OS 4.6.2.5](#)

[New in Avere OS 4.6.2.4.C1](#)

[New in Avere OS 4.6.2.3](#)

[New in Avere OS 4.6.2.2](#)

[New in Avere OS 4.6.2.1](#)

[New in Avere OS 4.6.1.4](#)

[New in Avere OS 4.6.1.3](#)

[New in Avere OS 4.6.1.2](#)

[New in Avere OS 4.6.1.1](#)

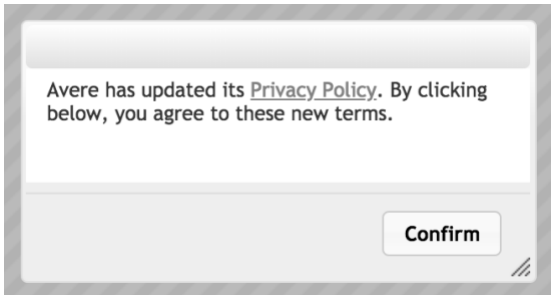
[Contact Support - Avere Global Services](#)

## New in Avere OS 4.6.2.5.C11

The Avere Privacy Policy has changed. Administrators should review the changed policy before using (or continuing to use) Avere OS support data upload features.

The current document can be read at <http://www.averesystems.com/privacy-policy>.

If you have not yet accepted the terms, a pop-up dialog appears when you load the **Support** tab or the **Cluster > Support** configuration page on the **Settings** tab of the Avere Control Panel. This dialog includes a link to the revised policy. Click the **Confirm** button to agree.



Links to the current terms of use and privacy policy statements have been added to the left margin of the **Support** tab.

## New in Avere OS 4.6.2.5.C10

This release updates Avere OS to support additional SSD drives now being used in Avere FXT 5000 Series hardware. (24852)

Note that if you receive any hardware that includes these drives, you must upgrade your Avere OS software to version 4.6.2.6.C10, and you cannot downgrade to a previous version.

Avere OS 4.7.3.2.C5 also supports these new components.

## New in Avere OS 4.6.2.5.C9

This critical release fixes the high-priority issue described below. This upgrade is appropriate for all customers running an earlier release of Avere OS version 4.6.2.5.

### Resolved Issue

#### Filesystem

23247 Fixed a race condition related to stale inodes and oplocks, which could cause a filesystem restart.

# New in Avere OS 4.6.2.5.C8

This critical release fixes the high-priority issue described below.

## Resolved Issue

### SMB/CIFS

- 23119 Code was changed to ensure that SMB client session setup operations for unknown users fail with a "no such user" error. Before this change, the system could spend resources attempting to establish the session before doing an access check. This change also helps prevent an out-of-memory issue due to repeated setup attempts from unknown users. (Share and filesystem ACLs are enforced regardless of this change.)

# New in Avere OS 4.6.2.5.C7

## Resolved Issues

### Cloud Core Filer Object Store (FlashCloud)

- 23049 Avere OS now allows new Google Cloud storage class names by default. Before this change, it was possible for core filer creation to fail for a supported Google Cloud storage type because its storage class name had changed.

# New in Avere OS 4.6.2.5.C6

## Resolved Issues

### Filesystem

- 22515 Changed procedures to ensure success returning ACLs for commands involving snapshot objects.
- 23176 Fixed an issue that caused the cluster to incorrectly identify an operation as stuck.

### SMB/CIFS

- 20126 Added a custom setting that Avere Global Services can use to help address LDAP server timeout issues.
- 23124 Additional work to correct an operation timeout when an SMB client accessed a path containing a relative symbolic link that would access content above the root of the Avere global namespace junction.
- 23638 Fixed an issue that caused poor performance in internal SMB path-name resolution when top-level parent directories were actively changing.

# New in Avere OS 4.6.2.5.C5

## Product Note

### Software Compatibility for SSD System Drives

Starting in December 2016, Avere FXT 5000 Series hardware nodes ship with solid-state (SSD) system drives instead of hard disk (HDD) system drives.

Support for SSD system drives in the 5000 Series was added in Avere OS version 4.6.2.5.C1. If your system includes FXT 5000 Series hardware nodes with SSD system drives, do not attempt to downgrade the cluster software to a version earlier than 4.6.2.5.C1. Downgrading these nodes to an incompatible software version causes hardware error messages with the phrase "Incorrect drive type installed, expected system drive."

This release of Avere OS (4.6.2.5.C5) includes a change to prevent nodes with SSD system drives from installing earlier, incompatible software.

## Resolved Issues

### Cloud Core Filer Object Store (FlashCloud)

- 22336 Fixed an issue in which an Avere filesystem check generated errors and did not complete in a reasonable amount of time.
- 22558 Fixed an issue where `rm -Rf` operations on cloud core filers could cause a filesystem restart loop.
- 22782 Reduced memory allocation overhead by pre-allocating buffer handles at boot time.
- 23046 Fixed a bug that could cause a filesystem restart when flushing files to cloud storage.
- 23276 Fixed an issue that could prevent the cluster from creating new cloud snapshots.
- 23360, 23369 Updated code to prevent a crash if filesystem request references an invalid cloud core filer ID. The system internally logs this situation now.

### Filesystem

- 21751 Fixed an issue that could cause node removal to stall while flushing zero-length files.
- 22546 Added code to improve success when flushing changed data to back-end storage; before this change, it was possible for write operations to cause a filesystem restart event.
- 22785 Fixed an issue where a null file handle pointer could cause a service restart.

### General

- 22258 Deleting subdirectories that had been involved in multiple data migrations could cause a panic because of an invalid (nlink) field count. This issue has been resolved.
- 22591 A system log file was moved to a different partition in order to avoid system crashes caused by adding a node or changing SNMP settings.
- 22624 Allow nodes to obtain valid configuration data when rejoining a cluster after an unclean restart. Before this fix, nodes with damaged configuration files would show as unjoined.
- 22688 Fixed an issue where a node reformat operation could cause a restart loop on the surviving nodes because the HA process was unable to converge.
- 22825 Updated the code for reporting power supply failure alerts to reduce the likelihood of false positives.
- 23281 Added an option to disable checking for unsupported SFP module cables.
- 23519 Software was changed to prevent downgrading FXT 5000 series nodes with SSD system drives to Avere OS versions below 4.6.2.5.C1. Earlier software versions do not support SSD system drives.

## NFS

- 23064 Reworked file lock monitoring code to address a problem where client facing IP addresses were not being correctly distributed among nodes.
- 23257 Changed ONC RPC alert procedures to log dropped connections every 60 seconds instead of raising alerts.
- 23383 Fixed an issue that could cause the ACL cache to become stale after a service restart if a pending *set* *FACL* record still needed to be flushed.
- 23398 A problem was fixed that could cause services to restart (and sometimes lead to a core dump) when core filer IP addresses were added or changed during NLM lock/unlock operations from clients.
- 23469 Improved memory allocation for NFS/SMB connector. Tests of this change have shown this module's memory consumption to be reduced by a factor of 10.

## SMB/CIFS

- 22953 Fixed a bug that could cause failures when rapidly opening and closing files on an ACL-enabled SMB share.
- 23139 Improved handling of responses to SMB client eventlog pipes, which are unsupported. Such queries now give an error response; before this change, they could cause an SMB service restart.
- 23197 Fixed an issue that sometimes caused the SMB service to restart when several concurrent Kerberos sessions were open by the same user.
- 23288 Added code stack trace options to assist identifying sources of memory fragmentation that could lead to SMB service restarts.
- 23339 Fixed a problem that could, in rare cases, lead to a filesystem restart when modifying the default ACLs at a junction that links to a cloud core filer.
- 23376 Added an optional process restart and core file if the SMB service detects an inconsistent oplock state during a file close operation. The core file will be used to help debug the source of the inconsistent oplock state, which can cause a process restart later when the file is opened.

# New in Avere OS 4.6.2.5.C4

## Resolved Issues

### Cloud Core Filer Object Store (FlashCloud)

- 22314 Prevent spurious "Bucket does not exist or is not accessible" warning messages for cloud core filers during a normal startup.
- 21555 Improvements to ensure point-in-time data consistency during cloud snapshots.
- 23110 Fixed a problem that could cause the cluster to ignore the cloud metadata database and attempt to read in a non-existent object from the bucket. The error could cause valid file metadata to be reported as stale.

### Filesystem

- 22894 Allow data to be discarded from the cluster cache after changes on the backend core filer have altered the root export filehandle or other information. Before this change, the FXT cluster was unable to flush some data back to the core filer; now, it logs the affected files' information and then discards the changes.



## SMB/CIFS

- 22826 Added support for Microsoft Active Directory Resource SID compression to avoid SMB client access errors when authenticating with Kerberos.
- 22915 Corrected an invalid DNS name in the service principal name used by a Kerberos cross-domain protocol transition ticket request.
- 22566 Improved the latency of SMB read and write operations by prioritizing the transmission of read and write responses above the admittance of new client requests.
- 22895 Added support to enable Avere cache-to-cache processing of SMB file and directory create operations that use Native Identity.
- 22922 Corrected a performance problem due to SMB find operations not being forwarded correctly during Avere cache-to-cache processing.
- 23124 Corrected an operation timeout when an SMB client accessed a path containing a relative symbolic link that would access content above the root of the Avere global namespace junction.

## vFXT

- 22817 Reworked memory handling code to correctly allocate NVRAM equivalent memory on virtual nodes. Before this change, the system could restart with an “out of NVRAM” error on virtual nodes, which do not have NVRAM.

# New in Avere OS 4.6.2.5.C3

## Resolved Issues

### Filesystem

- 22835 Fixed a memory overrun in the filename lookup path during snapshot file access.
- 22477 Fixed a race condition observed on systems configured for a cache policy of "Clients Bypassing the Cluster" (or similar). Before this change, stale file attributes could be cached if racing *write* and *get attributes* operations were issued for a file not in the cache.

# New in Avere OS 4.6.2.5.C2

The .C2 release includes one change, which allows Avere OS to recognize Google Compute Engine vFXTs after Google changed the string returned in response to an SMBIOS query. (22847)

# New in Avere OS 4.6.2.5.C1

## Resolved Issues

### Cloud Core Filer Object Store (FlashCloud)

- 22558 Fixed a problem that could lead to a filesystem restart when deleting large quantities of files on a cloud core filer. This change reduces memory resource contention by leaving some of the memory used for filesystem metadata unpinned.

## Filesystem

- 21771 Fixed a possible lock-order reversal when replaying file or directory metadata updates after an FXT node failover event. Before this fix, this issue could cause a filesystem restart.
- 22477 Fixed a race condition that could result in stale file attributes being incorrectly cached. This issue affected systems set for read caching and core filer verification, and was caused by a conflict between a write operation (passed through to the core filer) and a simultaneous getattr operation.
- 22565 Fixed an NFS readdirplus-related crash where certain cluster configurations with "always\_forward" enabled and certain conditions (the "HA Barrier" being in place,) would result in a filesystem restart.

## General

- 22708 Platform support for Revision E of the FXT 5000 Series hardware platform (updated system drive support).

## NFS

- 8217 Added cluster-wide and per-node client-connection counters to `nfs_front_rpc_vserver<ID>` statistics collection, and added the API call `stats.clientCounts` to retrieve them.
- 22738 Improved detection of NetApp and Isilon core filers to automatically populate the default Filer Class attribute (introduced in Avere OS 4.6.2.5).

## SMB/CIFS

- 21790 Fixed a use-after-free bug that caused sporadic filesystem restarts while processing ACL fetches from SMB/CIFS clients.
- 22516 A memory leak was eliminated in the handling of compound chained operations where the first operation in the chain is not an `smb2_create` operation.
- 22563 Fixed a statistics counting issue that caused double counts in `smb2_write`. The issue caused `smb2_write` replies to be double the call count, histogram bucket counts were double the responses, and error and time stats were doubled.
- 22717 Fixed a timeout issue that prevented access to an Avere-hosted SMB/CIFS share if the custom setting `vserver<ID>.NfsFrontEndSobuf` was configured to be higher than 1800000.
- 22743 Fixed a resource leak with the SMB/CIFS connection management database that resulted in occasional restarts of the SMB/CIFS process.

## vFXT

- 19543 Fixed an issue where virtual platforms had filesystem restarts due to crashes in the filesystem intent log management process.

# New in Avere OS 4.6.2.5

## 1. New Features and Enhancements

### Hourly Cloud Billing

The accounting mechanisms for vFXT nodes have been standardized across all supported cloud providers, and moved to a per-instance, per-hour model.

### Google Subnetwork Support

This release includes full support for GCE subnetworks. (The legacy network workaround described for a [previous release](#) is no longer necessary.) (21565)

### Improved Support for Eventually Consistent Cloud Storage

This release includes enhancements that improve communication with eventually consistent private object cloud filers, like HGST (Amplidata) and Cleversafe storage.

These changes prevent the situation where PUT calls could be retransmitted and cause a race condition within the cloud core filer's storage stack. (21937)

### Microsoft Azure Performance Improvements

Modifications were made to the Avere OS kernel for more effective use of the CPU and disk I/O resources in Microsoft Azure. vFXT systems hosted on Microsoft Azure now show better datapath performance.

### Performance Improvements for Windows ACLs

A more effective cache system for Windows ACLs is now being used in Avere OS, giving better performance for setting ACLs, retrieving ACLs, and processing new files that have ACL inheritance enabled.

Before this change, performance could be poor when adding ACLs and ACEs to small files.

### SMB %S Home Directory Share Support

Home directory dynamic shares in SMB now can be defined with %S, in addition to the previously available %U shares. Share-level ACLs cannot be applied to %S shares. (20246)

### Vendor Information Added for NAS Core Filers

Avere OS now stores information about the hardware vendor for NAS core filers. This change allows vendor-specific behaviors to be detected and handled.

Vendor information is stored in a new variable, `filerClass`, which is an optional argument in core filer XML-RPC calls. To set the vendor information from Avere Control Panel, use the new **Filer Class** field on the **Core Filer Details** page. The **Add New Core Filer Wizard** also includes a **Filer Class** selector in the first page for NAS core filers.

When you update a cluster to Avere OS 4.6.2.5, the system attempts to assign an appropriate filer class value for each core filer in the system. You can use the **Manage Core Filers** page to check the assignments for each core filer. (21406)

## 2. Resolved Issues

### Cloud Core Filer Object Store (FlashCloud)

21571	Improved the thread dispatching of X.509 certificate validation connections to prevent filesystem restart under certain error conditions
-------	--

- 21889 Improved the memory efficiency of the FlashCloud object cache while processing a large number of file deletions
- 22216 In some situations, canceling a long-running FlashCloud snapshot job could leave the system in synchronous write-through mode indefinitely, impairing performance. This problem has been corrected.
- 22233 Improvements were made to the file system's use of the OpenSSL API to better leverage thread safety. This prevents filesystem process restarts in high-rate SSL connection error-handling situations.

## Filesystem

- 20761 Fixed a race condition that could cause a filesystem restart if parallel rename calls for the same directory arrived on different nodes
- 20910 Fixed an issue where clients were not able to see the *.snapshot* entry from a NAS core filer with Local Directories enabled
- 21448 Fixed a race condition where certain types of directory updates from the directory manager were getting trapped in the system during random HA-barrier events
- 21762 Improved the handling and subsequent removal of incorrectly configured NAS core filers. This situation affected core filers with Local Directories enabled. If these core filers were not configured properly for access by the Avere cluster, they could not be populated. In this case, the dashboard showed the condition "Not making progress in populating the directory."
- 21795 Fixed a timing issue that could cause a "stuck revoke" condition to appear in the Dashboard. The issue involved writeback flushes to newly created and modified files.
- 21963 Dashboard alerts were improved to suppress repetitive "error=IO" alerts, which could flood the dashboard and the configuration database
- 21998 Fixed a filesystem restart caused by a race condition in cookie cache pruning and other maintenance tasks
- 22005 During a cluster-wide HA failover, it was possible for the state of file attributes in a mirrored copy to block client access to these files. This problem has been fixed.
- 22038 Fixed an issue that could cause restarts when a file in the Avere cache was also modified directly on the back-end storage system. This issue occurred in a read-only cache mode when the back-end file size was smaller than the file size expected by the Avere cache. The size mismatch caused errors when the Avere cluster passed a client's write request through to the core filer.
- 22039 Fixed a null pointer dereference when reading from the cache under certain error conditions. This problem could cause a filesystem restart.
- 22114 Fixed a race condition that involved parallel FSSTAT NFS calls received at multiple FXT nodes for an uncached file handle
- 22139 Improved the flushing efficiency when writing large files to cloud core filers to allow the FlashCloud snapshot to complete more quickly
- 22174 Fixed an issue where file-extending write calls from NFS clients with the UNSTABLE flag could lead to improperly zeroing out parts of 16KB ranges
- 22193 Improved the response time of NFS 'setattr' calls against files that are actively being flushed to the core filer
- 22230 Enhanced performance by improving response times for client NFS writes using the UNSTABLE parameter
- 22241 Improved performance of client writes (with the UNSTABLE flag) to large files that are simultaneously being flushed to the cloud core filer
- 22312 Fixed a filesystem restart issue caused by the use of the non-default "firstSegmentBytes" custom setting
- 22335 Fixed a filesystem restart issue in which parallel directory removal calls could cause resource depletion
- 22348 Fixed an issue where the Dashboard could be flooded with repetitive alerts for failures to populate the persistent directory cache with metadata from a NAS core filer

22360 Fixed a filesystem restart issue that was triggered over time by clients browsing *.snapshot* data on a NAS core filer

## General

- 16564 Added a new option to the cluster setup screen for unconfigured nodes. Click on **I want to update the system image first**, and below the **Get New Image** section there is a **Upload New Image** section.
- 21176 Updated the Add New Core Filer wizard to ensure that the **Certificate verification** selection box displays **OCSP\_CRL** when Google Cloud Storage S3 is chosen as the service type
- 21885 Improved the Dashboard alerting and internal handling of errors from a NAS core filer when certain limitations (maximum number of files in a certain directory) are reached on the NAS core filer
- 21931 Fixed the ordering of "Core filer network name/IP" entries in the Avere Control Panel so that they are not reordered
- 21947 The use of the # character is now forbidden in core filer and vserver names
- 22309 Fixed a resource deadlock in the cluster control plane communications path. This problem was observed when adding new FXT nodes to an existing cluster with more than 22 nodes, and could cause the cluster join process to stall.

## NFS

- 20960 Fixed a problem that caused assigned export policies to be set back to the default policy after changing a NAS core filer configuration to Local Directories mode

## SMB/CIFS

- 18053 Fixed a rare condition that could cause an SMB restart under highly concurrent client access workloads across a common dataset. This problem only occurred when internal garbage collection processes were also running.
- 19934 Removed remaining unused registers associated with the removed serverid.tdb database that could cause spurious watchdog reports
- 20718 Enhanced the performance of f "setfacl" calls (creation of Windows ACL metadata by clients) when using FlashCloud core filers. This improves robocopy performance when ACLs are being copied.
- 20742 Fixed a resource management race condition that could cause SMB to restart after simultaneous client connects/disconnects
- 21544 Fixed an issue with the NFS-SMB connector that could cause the SMB process to restart
- 21651 Improved the Previous Versions snapshot browsing experience for SMB users and eliminated some related SMB restarts
- 21837 Improved the stats.hotClient() API call to include and differentiate information from NFS and SMB/CIFS clients. This information also appears on the Dashboard in the Clients tab.
- 21870 Fixed a resource leak in the SMB process when handling large quantities of outstanding connection requests
- 21997 Fixed a filesystem restart issue that was triggered during Windows ACL migration by a FlashMirror job that encountered errors and required rollback
- 22126 Improved the memory management of Windows ACL information for files that have a large number of ACEs in the ACL; this issue could cause SMB process restarts
- 22151 Fixed an SMB process restart that was observed during near-simultaneous opening of the same file across multiple FXT/vFXT nodes
- 22157 Improved support for base64-encoded canonical names stored in LDAP being queried by Avere Directory Services

- 22158 Fixed a process restart in the SMB protocol stack that was triggered by client-initiated directory creation with ACLs against a Native Identity SMB share
- 22286 Improved TCP window sizing to enhance performance of client SMB connections

## vFXT

- 21968 Fixed an issue that caused the AWS/GCE Cluster Manager to fail during the software package installation on new vFXT nodes
- 21794 Enabled Azure vFXT nodes to read SSH keys from the Azure provisioning environment
- 21975 Fixes the previously documented [issue](#) that prevented new software images from being downloaded to vFXT nodes on a private network with a proxy server
- 21981 Added support for 375GB local SSD storage (for GCE) and 8TB persistent cache storage (for all cloud platforms) when creating a vFXT cluster
- 22111 Fixed a defect that prevented cluster destroy actions when the Avere OS Cluster Manager was used in AWS regions other than us-west-2.

# New in Avere OS 4.6.2.4.C1

## 1. New Features and Enhancements

### Microsoft Azure vFXT Early Access

Avere OS 4.6.2.4 contains early-access support for Microsoft Azure DS13 virtual machine vFXT nodes. Contact your Avere sales representative to learn how to participate in an early access program.

### Streamlined Deployment Process for vFXT Nodes

The processes used for creating virtual FXT nodes have been consolidated, and the steps to create vFXT clusters and vFXT node instances now are more consistent across all supported cloud computing providers.

Before this change, Amazon Web Services EC2 vFXT clusters were configured using CloudFormation templates. Now, instead of templates, AWS vFXTs use a virtual machine instance called the Cluster Manager to create the vFXT clusters.

Previously created vFXT nodes and clusters are still valid; this change only affects new node or new cluster creation.

vFXTs created on Google Compute Engine products have always used a Cluster Manager VM (referred to in some earlier documents as the “shepherd”) to create vFXT nodes. Microsoft Azure vFXTs also use this system.

To learn more about creating virtual FXT nodes, read the vFXT installation guide for your cloud platform:

- [vFXT Installation Guide for Amazon Web Services](#)
- [vFXT Installation Guide for Google Cloud Platform](#)
- vFXT Installation Guide for Microsoft Azure - *Available to early access customers only; contact your Avere sales representative or Avere Global Services to learn more.*

Additional tips about this new deployment process are included in the next section under [vFXT Guidelines](#).

### vFXT Performance Enhancements

#### Single Root IOV Support

This release adds support for the single root I/O virtualization (SR-IOV) driver, which is supported for vFXTs on Amazon EC2. This driver improves network performance and reduces latency when running r3.8xlarge vFXT instances.

Note that this driver is unsupported on older versions of Avere OS, so you must ensure that all nodes in your cluster have updated software before enabling the driver. Follow the procedure in [Enabling SR-IOV on Amazon vFXT Clusters](#) to turn on this feature. (21666, 17755)

## CPU Scheduling Improvement

Avere OS performance was enhanced on large core-count AWS and GCE instance types by improving CPU scheduling.

## 2. Product Tips and Workarounds

This section documents known issues and best practices when using Avere OS 4.6.2.4.

### Do Not Use “Dot” Characters In AWS Buckets for Cloud Core Filers

When creating buckets hosted on Amazon Web Services (AWS) S3 for use as core filers, do not include a period (.) in the bucket name. Avere software does not allow periods in cloud core filer bucket names. (20422)

### Changing Node Naming Setting Does Not Update Existing Node Names

A change to the node name settings in the Avere Control Panel’s **Cluster > General Setup** settings page currently does not change the names of existing FXT nodes in the cluster. However, new nodes added to the cluster after customizing the node name settings (node prefix and first node number) do receive names based on these settings.

Individual node names can be changed on the **Node Details** page. Click a node name from the **Cluster > FXT Nodes** settings page to load this page.

Avere is investigating this issue. (21439)

### Unsupported Atmos Cloud Storage Option Appears in Control Panel

The Avere Control Panel wizard for creating cloud core filers currently includes an option to select EMC Atmos cloud storage as the cloud service provider, but EMC Atmos storage currently is not supported for use with Avere clusters. The option exists to support testing and early access users as Avere validates Atmos for use as a cloud core filer.

If you have questions about EMC Atmos early access, please contact Avere Global Services. (21646)

### Use HTTPS To Enable Software Update Through a Proxy

When downloading a new Avere OS software package on a cluster that uses a proxy server, you must use secure HTTP (https://) when specifying the download URL. (This feature is available on the **Software Update** settings page of the Avere Control Panel.) A regular HTTP connection for software download will not be routed through the proxy server. (22054)

### Proxy Must Be Manually Selected When Creating a New Cloud Core Filer

If you add a new cloud-based core filer to a cluster that uses a web proxy, the proxy is not automatically used for communication with the new core filer. To use the proxy for core filer communication, select it manually in the third page of the **Add New Core Filer** wizard. The third page includes settings for service type, cloud credential, compression, and other options; use the **Proxy** selector at the bottom of the panel to choose the proxy configuration for communication with this core filer.

There is no way to add the proxy after defining the core filer; instead, you must remove the incorrect core filer and add it with the proxy configuration included.



## vFXT Guidelines

The following tips apply to vFXT clusters in cloud computing environments.

### Supported vFXT Instance Types

The currently supported instance types and per-node cache sizes are shown in the table below. Cache sizes are not restricted by instance type; any combination of supported instance type and cache size is valid.

	Cluster Manager Instance Types	vFXT Instance Types	Supported Cache Sizes
<b>AWS</b>	t2.large	r3.2xlarge r3.8xlarge	1TB EBS, 4TB EBS
<b>GCP</b>	n1-standard-4	n1-highmem-8 n1-highmem-32	1TB P-SSD, 4TB P-SSD, 1.5TB Local SSD

### Cluster Manager Can Show Deleted Clusters

If a vFXT cluster is destroyed but is not explicitly deleted from the Cluster Manager dashboard, the dashboard might continue to show information about the cluster, even though the information is not accurate. Also, if the management IP address from the deleted cluster is assigned to a new cluster, the dashboard will show two clusters with the same management IP address and the old cluster's (invalid) information.

To avoid this confusion, explicitly delete the cluster from the Cluster Manager list when removing it.

Avere is reworking this code to resolve the issue in an upcoming release. (21893)

### vFXT Cluster Manager Passwords Must Not Start or End with Spaces

Because of a recent change, passwords with leading or trailing spaces are not correctly stored for the Cluster Manager. Until this issue is resolved, do not use blank spaces at the beginning or end of passwords for vFXT clusters. (21852)

### Check for Correct Privileges Before Creating Clusters or Cluster Managers

Avere software currently does not verify the user's cloud credentials before attempting to create a vFXT cluster. You must make sure that appropriate privileges exist to instantiate and configure VMs on the cloud service being used. (21660)

It is possible to instantiate a Cluster Manager instance that does not have sufficient privileges to create vFXTs. The error messages that Avere software returns in these cases are not specific enough to suggest the root problem. Avere is working to remedy this situation.

In AWS systems, this type of error typically is related to either a missing Identity and Access Management (IAM) role or an invalid IAM policy. The solution depends on the error:

- A missing IAM role cannot be added after the Cluster Manager instance has been created. Destroy the instance and create a new instance with the correct role.
- If the Cloud Manager instance was created with the correct role but the IAM policy prevents it from working, the policy usually can be modified.

To learn about IAM roles and policies, read the Amazon EC2 documentation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>. (21946)

### vFXT "Power Up" Button Does Not Work

A button appears in the Cluster Manager to start a vFXT instance, but this button currently fails to power up the vFXT. Use the cloud provider's start controls to power up a node. (20295)

### Ignore IPMI Error Message when Renaming a vFXT Node

If you rename a vFXT node (by editing it in the **Node Details** page of the Avere Control Panel), an error message appears about modifying the IPMI mode. This error can be safely ignored. The message reads "Could not modify the IPMI mode of the node due to the following error: Permission denied. Please try again," however, the name has been successfully changed and a retry is unnecessary. (22077)



## Non-AWS Credentials Cannot Be Selected in Avere Control Panel

vFXT clusters have a setting on the **Cluster > Administrative Network** page for selecting the cloud networking credential, but currently this page recognizes cloud credentials from Amazon only. Cloud credentials associated with other providers do not appear in the options list and currently cannot be selected by using the Avere Control Panel.

The workaround is to use the Avere OS command-line interface to set the cloud credential for a GCE or Azure vFXT cluster. Contact Avere Global Services if you need additional information.

Non-AWS cloud credentials do appear in the **Cloud Credentials** list, but on this page they can only be modified or removed. (22050)

## Ensure Proxy Connectivity During vFXT Cluster Upgrade

When upgrading an existing vFXT cluster that uses a web proxy to Avere OS 4.6.2.4, the proxy is automatically used to access cloud provider API endpoints during the upgrade. If the proxy server is unable to reach the endpoints, the upgrade might stall. In that case, remove the proxy configuration so that the upgrade can proceed.

Note that when specifying a proxy URL, you must include the `http://` or `https://` prefix. Neither the Avere Control Panel nor the command-line equivalent enforce this requirement. (22079)

## Invalid Proxy Configuration on a Private Network Can Make Cluster Manager Unreachable

If you create a Cluster Manager on a private network that routes traffic to cloud endpoints through a proxy server, the Cluster Manager must always have access to a valid proxy configuration. If there is no valid proxy, the Cluster Manager can become unusable.

Extra care should be taken to avoid this issue, including the following:

- On first login to a newly created Cluster Manager, be prepared to correctly enter a valid proxy configuration at the initial prompt.
- Do not remove a proxy configuration from the Cluster Manager while the proxy server is required for access to the cloud endpoints.

If a Cluster Manager becomes unusable because of a bad proxy configuration, recreate the Cluster Manager and destroy the old instance. (22072)

## Google (GCP/GCE) vFXT Issues

The notes in this section affect vFXT clusters hosted on Google's cloud service.

### ***GCE: Cannot Specify IPs when Adding New Cluster Nodes***

Although there is a "network override" feature available that allows you to choose specific IP addresses when creating a new vFXT cluster, there is no IP override feature available when adding a cluster node. However, Avere OS will always use an IP address assigned to the cluster if one exists; if not, it will assign a new IP address range.

To ensure that new nodes receive IP addresses in the desired range, make sure that extra IP addresses are available in the cluster.

1. Use the Avere Control Panel to assign enough cluster IP addresses to cover the new nodes.
2. Use the Cluster Manager to create and add the nodes.

(22059)

## Amazon (AWS EC2) vFXT Issues

The notes in this section affect vFXT clusters hosted on Amazon's cloud service.

### ***Unique CIDR Block Must Be Used for Each AWS Multi-AZ Cluster***

When creating multi-AZ clusters in AWS and allocating IP addresses by using CIDR, a unique block of addresses must be used for each cluster. If the same CIDR block is used for more than one AWS multi-AZ cluster, one IP address can be assigned multiple times (once in each cluster), leading to possible data corruption. Avere is investigating a solution for this issue. (21862)

## Enabling SR-IOV on Amazon vFXT Clusters

When using the single root I/O virtualization driver (new in this release), it is important to upgrade software on all cluster nodes before activating the driver.

Note that this feature is supported on AWS EC2 r3.8xlarge vFXT clusters only.

Follow this procedure to upgrade software and activate SR-IOV on clusters running Avere OS 4.6.2.3 or earlier software versions:

1. Use the Software Update settings page (or a command-line equivalent) to update the cluster software to Avere OS v 4.6.2.4 or higher.
2. Wait until all nodes in the cluster have updated and no alert conditions appear on the Avere Control Panel.
3. Verify that the software upgrade process has completed by checking the Alerts section on the Avere Control Panel Dashboard tab for a completion notification.

**CAUTION:** The verification step is *very important*. Attempting to enable SR-IOV on a cluster running Avere OS versions older than 4.6.2.4 will make the instance unusable and may result in loss of unflushed cache files.

4. Shut down all vFXT nodes in the cluster, using one or more of these methods:
  - Use the **FXT Nodes** page in the Avere Control Panel to power down each node
  - Use the **System Maintenance** page in the Avere Control Panel to power down the cluster (verify that each vFXT responds to this command before proceeding)
  - Use the AWS console to issue a “stop” command for each vFXT instance
5. Use the AWS or EC2 command line tools to enable SR-IOV support on each vFXT. You must run one of the following commands for each instance ID in the cluster.
  - If using the AWS command line tool:  
aws ec2 modify-instance-attribute --instance-id *INSTANCEID* --sriov-net-support simple
  - If using the EC2 command line tool:  
ec2-modify-instance-attribute *INSTANCEID* --sriov simple
6. Start the vFXT instances again by using the AWS console or command line.

## 3. Resolved Issues

### Cloud Core Filer Object Store (FlashCloud)

- |       |   |
|-------|---|
| 20422 | The dot, or period, character (.) cannot be used in an AWS bucket name that is used as a cloud core filer. The cloud core filer creation wizard has been updated to reject bucket names that contain a dot.                                   |
| 21454 | Fixed a race condition related to locks on internal data structures during system shutdown. This condition could cause problems when cloud core filer operations were being processed during shutdown.  |
| 21557 | Optimized identification of unmodified data in the cache during cloud snapshot creation. This change reduces the amount of time the system remains in write-through mode during snapshots, improving performance for active client workloads. |
| 21778 | Updated the Avere API documentation for the corefiler.createCredential and corefiler.modifyCredential API calls.  |

### Filesystem

- |       |   |
|-------|---|
| 21562 | Added a throttle for write-intensive workloads that have been forwarded from one FXT node to another in the cluster under resource pressure. Before this change, forwarded write calls could cause buffer exhaustion in the cache, leading to process restarts. |
| 21605 | Preserves customized cache policy definitions when certain older clusters are upgraded to Avere OS version 4.6.2. This issue only affected systems that had enabled and then disabled Cache Utilization Controls before upgrading to 4.6.2.                     |

- 21690 Resolved a race condition when a client requests small reads against a file that is not cached. An early response optimization used for this scenario caused process restarts and has been fixed in this release.
- 21693 Fixed a filesystem process crash caused by a race condition related to locking the parent directory during concurrent rename or remove calls.
- 21791 Improved memory buffer resource allocation for pre-fetched reads to prevent process restarts during heavy and highly parallel client read activity on files that are not currently in the memory cache.
- 21882 Fixed an issue that could cause the filesystem process to exit and restart if Avere OS filesystem diagnostic tracing encounters an incomplete operation.

## FlashMove/FlashMirror

- 21589 In data management jobs involving cloud storage and done between v4.6.1.4.C2 and v4.6.2.3, there is a chance that if the migration job encountered an error while moving data, data on the destination core filer could become misaligned with what resides on the source. The source data is not impacted in any way. An upgrade to v4.6.2.4 is recommended to eliminate any exposure.
- 21759 Fixed an issue that could cause clients to incorrectly reference the source files instead of the (incomplete) destination files after a FlashMove or FlashMirror job was aborted or failed. This situation occurred when a junction to the destination was created after the migration failed.
- 21733 Fixed a problem that prevented FlashMove/FlashMirror jobs from succeeding on EMC Isilon core filers. The problem occurred when migrating data that involved SMB/CIFS ACLs; the 'admin' user specified at job creation time was not being mapped to the root user for certain operations to the core filer.
- 21801 Improved client access by eliminating an unnecessary flush phase between the "Transitioning Clients" and "Full Caching" modes for FlashMove or Flash Mirror data migrations.

## General

- 15760 Changed wording in a software upgrade wizard dialog to clarify that uploading a new software image on the Avere cluster does not automatically activate the image. Before this change, the wording incorrectly suggested that loading the software also initiated the activation step.
- 16564 Added a new option to the new cluster configuration wizard to allow uploading a new system image as an alternative to downloading one from a URL. These options appear after clicking on the link "I want to update the system image first".
- 20983 The NfsExportTask process was assigned to an incorrect scheduler and priority, causing repeated "long-wait" dispatch error messages in the log. This task has been assigned to the proper scheduler.
- 21021 Fixed a bug in the alert notification email client code that prevented certain alerts from being properly sent via email.
- 21412 Fixed a bug that caused the statistics and counters history management process to restart after one year of up time.
- 21505 Modified kernel scheduling behavior to avoid a situation where interrupt threads could starve threads that were pinned to a particular CPU, even when other idle CPUs are available.
- 21552 Improved the statistics counters that report memory buffer resource pressure while the system is under heavy load.
- 21570 Increased system buffer resource limits to handle a higher quantity of in-flight filesystem operations.
- 21617 Fixed an issue that caused Internet Explorer version 11 to fail at uploading cloud encryption recovery keys to an Avere cluster running v4.6.2.
- 22084 Fixed a problem preventing the multi-cluster dashboard from loading that was introduced in 4.6.2.4

## Security

- 21576 Upgraded OpenSSL to version 1.0.2g to include high-severity fixes and remove DROWN attack vulnerability. Note that Avere OS uses TLS instead of SSL; SSL v3.0 is now explicitly disabled.

## New in Avere OS 4.6.2.3

### 1. New Features and Enhancements

#### API Updates for Amazon Web Services Compatibility

The API for configuring AWS clients was updated to use Amazon Signature v4 and to correctly find the CA certificates file after software upgrades. This change improves usability when working with updated Amazon endpoints or connecting in an isolated AWS region. (21375)

#### Workaround: Use Legacy (non-subnet) networks for Google Compute Engine vFXT Clusters

A change in the default Google Compute Engine (GCE) network settings has introduced an incompatibility with Avere vFXT cluster instances.

Avere Systems is in the process of adding support for GCE subnetworks, but for now users must create a “legacy” network for their new GCE project. This setting overrides the new GCE default and creates a system with a single subnet. Read <https://cloud.google.com/compute/docs/subnetworks> and [https://cloud.google.com/compute/docs/networking#creating\\_a\\_legacy\\_network](https://cloud.google.com/compute/docs/networking#creating_a_legacy_network) to learn more.

For more information or help with configuration, contact Avere Global Services.

### 2. Resolved Issues

#### Cloud Core Filer Object Store (FlashCloud)

- 20755 A new argument was added to the `corefiler.createCloudFiler` API command to check for existing FlashCloud object data when re-creating a cloud core filer backed by a cloud bucket.
- 21459 Implemented a more dynamic write-cache limit for cloud core filers in the hour before a snapshot. This change increases efficiency of the write-back cache and reduces unnecessary write-back activity.
- 21488 Fixed a filesystem service restart that occurred when file rename operations were flushed to the cloud core filer under certain resource conditions.
- 21535 Adjusted early flushing of write-back data to cloud core filers when preparing for scheduled snapshot synchronization to prevent immediate flushing of recently-written data.

#### Filesystem

- 20905 Improved data placement efficiency to speed performance when flushing a file that has been modified on multiple cluster nodes.
- 21256 Improved handling of error conditions to prevent filesystem restarts during transition between source and destination in FlashMove and FlashMirror operations.
- 21383 Fixed a problem listing snapshot directories on an HNAS filer from NFS clients when Local Directories caching is enabled.
- 21414 Improved prioritization of non-empty files when flushing data to the core filer to free cache space.
- 21428 Fixed a race condition that could cause a filesystem restart when a truncating-setattr operation was attempted against a file that was being flushed to the core filer.
- 21492 Improved handling of larger (512KB) `readdir` and `readdirplus` calls to Oracle ZFS core filers. Before this change, filesystem service restarts were observed in v4.6.2.2 when using ZFS core filers.

## General

- 18670 Fixed a problem that prevented changes to static IP routes from taking effect, although the change was acknowledged. This problem affected changes from both the Avere Control Panel and the API.
- 21180 Improved the resilience of clusters after a node failure. Prevents a rare situation where a failed node could not be removed from the cluster and a new node could not be added.
- 21323 Fixed issues with resynchronizing cluster-wide configuration database information when the database is large and growing (the “starvationCheckInterval” assertion). These issues caused restarts of the filesystem service, which could possibly destabilize the cluster.
- 21405 Corrected the output of the 'corefiler.get' API call, which was incorrectly listing the same IP address(es) more than once. This was seen in systems with multiple core filer network ranges.
- 21497 Fixed an issue that resulted in the Multi-Cluster Dashboard being unable to locate the configurations of previously monitored clusters after upgrading software from v4.6.1.x to v4.6.2.2.
- 21529 Before this change, the Multi-Cluster Dashboard (MCD) process would not display any monitored cluster status or statistics if the number of monitored clusters exceeded 10 clusters. The limit is now 200 clusters per MCD instance.
- 21547 Fixed a regression in the stats.hotFiles XML RPC call that caused it to fail with the error <Fault 100: "Internal error: 'ops'"> when called with more than one optype. This fix also addresses the issue that caused the hot files dashboard to show all counter values as zero.

## SMB/CIFS

- 20882 Improved the ACL process that translates SIDs to user IDs. Before this change, slow responses from AD queries could cause ACL process restarts.
- 21365 Fixed an error that could cause ACL server communication problems after a “mirror reverse” request during a SMB ACL FlashMirror job. This issue could lead to interrupted SMB/CIFS client connections.
- 21411 Improved the process for parsing timestamps embedded in snapshot filenames when they are displayed in the Windows “Previous Versions” tab.

# New in Avere OS 4.6.2.2

## 1. New Features and Enhancements

### Junction-Level Export Policies for Cloud Storage

A new feature allows users to set export policies (which control client access to the files on GNS-exported core filer volumes) at the junction level in addition to at the core filer export level. This feature gives administrators more flexibility when working with core filers that export a single root directory (for example, “/”). Rather than applying a single export policy to the entire core filer, administrators can use this feature to assign different policies to virtual subdirectories (junctioned within the cluster’s global namespace) on the core filer. (16075)

Setting junction-level policies is optional; if no policy is set, the junction inherits the policy set on its core filer export and the behavior is equivalent to Avere OS 4.6.2.1 and earlier.

An explicitly set junction-level policy will always have precedence over an inherited export-level policy.

This feature can be used with FlashCloud core filer storage only. Creating virtual subdirectories as GNS junctions on a FlashCloud core filer is a way to impose a more granular access control on the FlashCloud core filer without having to directly connect to the cloud export and create each subdirectory path from an NFS client. However, limitations apply, and character encodings could have unexpected consequences - contact Avere Global Services for updated information about using this feature to create a subdirectory filesystem structure with unique export policies.

## Setting Junction-Level Export Policies

Export policies are configured in the **Settings** section of the Avere Control Panel.

Under **VServer**, open the **Namespace** page to show existing junctions. When adding or editing a junction, a new **Export Policy** option appears that you can use to select the policy.

Namespace -- gns

The screenshot shows a web form titled "Modify Junction /amzstd-1-tub-shadow". At the top right is a button labeled "Add New Junction". The form contains the following fields and controls:

- "Core filer admin name": A dropdown menu with "amzstd-1" selected.
- "Core filer export": A dropdown menu with "/" selected.
- "Export subdirectory": An empty text input field.
- "Auto-create subdirectory": An unchecked checkbox.
- "Export policy": A dropdown menu with "Inherit (default)" selected. An orange arrow points to this dropdown.
- "Advanced": An unchecked checkbox.

At the bottom right of the form is a "Submit" button.

Use the **Export Rules** page to create or edit export policies.

A best practice is to create individual junctions for subdirectories on the core filer, and avoid exposing the root of your core filer export. Subdirectory junctions give you the flexibility to customize access for each sub-tree on the core filer.

If you do need access to the core filer root path, either remove the junction pointing to the root or set a more restrictive export policy for the root than is used for the junctions pointing to its subdirectories. A client with access to the root has de facto access to any subdirectory beneath it in the core filer filesystem, so the root junction's export policy should be the most restrictive policy assigned to the core filer.

As always, removing a junction does not affect the file structure on the core filer, and does not remove any files in back-end storage. Removing a junction only removes the Global Name Space access to that path.

## Upgrading from Earlier Avere Software

When an Avere cluster upgrades to Avere OS 4.6.2.2, junctions retain their inherited export policies. Use the **Namespace** page (shown above) to add junction-level policies after the OS update.

If you have an existing root-level junction, a best practice is to replace it with subdirectory junctions. After creating the subdirectory junctions, either remove the root junction, or configure its export policy to be the most restrictive policy in the system.

## FXT 5000 Series Hardware Availability

This release provides full support for the new Avere FXT 5400 and FXT 5600 Edge Filers. The FXT 5000 Series delivers more processing power, expanded DRAM and NVRAM, significantly more SSD storage, and double the number of 10GbE data ports per node when compared to earlier Avere Edge Filers.

## Using FXT 5000 Series Hardware with Other Hardware

Because the FXT 5400 and FXT 5600 have faster processors and much more cache storage space than older Avere Edge Filers, Avere recommends using 5000 series Edge Filers with other 5000 series hardware to take full advantage of the new systems.

When used in a mixed cluster with less powerful Avere Edge Filers (including models from the 4000 or 3000 series), a 5000 series Edge Filer can be restricted by the cache size and speed of the smaller nodes. When an FXT 5000 is used in an HA pair with an FXT 4000 series or FXT 3000 series node, the 5000 series node will not use its entire cache quota, and data will be flushed from the cache earlier than it would be if used with 5000 series nodes exclusively.

Contact Avere Global Services to learn more about incorporating the FXT 5000 series into your environment.

## Read/Write Block Sizes Increased

This release adds an option to increase the maximum rsize/wsize settings for client-cluster communications to 512KB, from 64KB. Changing the size yields significant speed improvements, especially when reading and writing larger files. (20044, 20116)

The maximum block size for backend writes (moving cache content back to core storage) was also increased, to 1MB. This change gives a significant speed improvement for cloud storage and other systems that can accept larger write sizes. (20088)

Both of these options are configured through API settings only. Contact Avere Global Services to learn more about changing these settings.

## VServer Read-Only Optimization Mode for CIFS

This release adds the ability to set a vservers in “read-only optimized” mode to improve efficiency when handling CIFS requests for read-only volumes. The feature can be enabled by using the API or from the **CIFS** page in the **VServer** settings section of the Avere Control Panel. (20268)

When enabling this mode, you must stop and restart CIFS before the change takes effect. Follow these steps:

1. Disable CIFS on the vservers
2. Enable read-only optimization for the vservers
  - a. If using the Avere Control Panel, also re-enable CIFS in this step.
  - b. If using the command-line API, you must re-enable CIFS in a separate command after turning on optimization.
3. Wait for all conditions to clear

## Core Filer Available Space Display Clarified

The Avere Control Panel dashboard tab that shows core filer capacity now uses raw numbers instead of percentages in the table of available read and write space per core filer. Also, read-only core filers now show “-” for all values related to write space.

## NTFS ACL Support for NAS Core Filers not Connected to the Default Cluster Network

A change was made to enable the Avere cluster to properly source TCP connections by using Avere IPs that are able to reach the specific core filer IP networks. The Avere OS process (smbacl) that reads and writes ACL information from a NAS core filer’s SMB/CIFS server now can connect to core filer IP addresses that are not routable from the Avere default cluster network. (21326)

# 2. Resolved Issues

## Cloud Core Filer Object Store (FlashCloud)

20388	Fix for a condition that could cause the filesystem service to crash or restart when removing nodes from the cluster or rebalancing directory managers
20954	Optimized the number of keys requested per LIST call for cloud core filers
21238	Fix for filesystem service restart caused by race condition when recycling object cache memory
21241	Prevent the dashboard from reporting a cloud core filer connectivity condition that is no longer present
21255	Fix for filesystem service restart caused by race condition when handling same-directory update operations
20179	Fix the slowdown experienced when frequent updates to a single directory cause decreased throughput. Rsync is one example of a tool that exhibits this behavior. Avere recommends using the “--in-place” option for rsync jobs run against FlashCloud.



- 20227 Reduce the read overhead for overwrites and random writes flushed out to cloud storage
- 20299 Fix for FlashCloud optimized read-ahead memory leak that resulted in a filesystem service restart
- 20829 Improve error handling for object storage metadata responses observed during failovers and restarts
- 20841 Increase the default per-directory maximum subdirectory limit from 32,000 to 100,000
- 20954 Improve handling of LIST responses containing more than 16KB of data
- 20972 Fix a link reference count issue that could cause a persistent stuck flush operation condition
- 21224 Correct a cloud object version tracking issue that was causing filesystem service restarts
- 21245 Improve the handling of cloud read-ahead worker threads that encounter internal errors

## Filesystem

- 19580, 20530 Improve the accuracy of junction and export path names reported in alerts and hotfiles
- 20741 Improve convergence time of “Scheduled write-through” when large files are continuously being updated
- 21131 Fix for crash/restart due to block allocation deadlock
- 21133 Improve buffer reservation to support file updates in large directories
- 21190 Fix for filesystem restart during log replay when recovering from an HA failover event
- 21233 Improve error handling of ERR\_IO received from an NFS core filer during directory metadata flush

## General

- 16738 Provide a “Forceful Remove” button to force a stuck core filer removal operation to complete
- 17280 Improve filtering and display of IP address conflicts raised as Dashboard Conditions/Alerts
- 19322 Ability to send syslog and SMTP messages to an admin-specified TCP/UDP port
- 21075, 20823 Fixes a recently introduced bug in the alert email test available from the Avere Control Panel
- 21017 Suppress messages log chatter related to “vcm: getName”
- 21180 Implement an interlock to prevent simultaneous addition/removal of nodes in the cluster
- 21201 Fix the API xml-rpc corefiler.create command and its help text to handle reworked Cache Policy selection
- 20245 Improved node replacement workflow in a multi-zone Amazon Web Services EC2 cluster
- 21258 Fixes a bug that prevented setting a custom value for Directory Services polling

## NFS

- 20353 Improves resource management for UDP NULL RPC operations received from NFS clients using autofs
- 20611 Fixes a bug that caused a filesystem service restart after a failed UDP mount operation
- 21028 Adds support for large ONC/RPC MOUNT responses received from the core filer back-end storage

## SMB/CIFS

- 19871 Eliminate superfluous messages in the SMB/CIFS logging component
- 20176 Bug fix for NFS setattr that fail because of SMB file ownership semantics
- 20743 Fix SMB “find” queries without wildcards to properly resolve symbolic links
- 21208 Allow CIFS to be enabled when usernameSource is set to AD and native identity is disabled



# New in Avere OS 4.6.2.1

## 1. New Features and Enhancements

### Redesigned Cache Policy Settings

Significant improvements were made to the Avere cache policy options and to the Avere Control Panel pages for setting and modifying cache policies. The improved features make the settings easier to understand and easier to use, and help administrators avoid setting inappropriate cache policies.

Key features of this redesign include:

- **An expanded set of system default cache policies** - Five default policy options now exist, to cover most use cases. Users also can create custom policies.
- **Intuitive cache policy names** - Each of the five policies has a descriptive name based on use cases.
- **Diagrams and details on the cache selection page** - When selecting a policy, the Avere Control Panel pages show a system diagram and describe the cache settings, use cases, and system requirements.
- **Warnings about policy transition** - The default policy descriptions include warnings about steps that will be required to transition from this policy to another policy. These warnings help users decide which cache policy is best now and in the future.

Cache policy settings appear in the **Add New Core Filer** wizard and in the **Core Filer** section of the **Settings** tab on the Avere Control Panel.

### Default Cache Policies

The table below describes the five default cache policies and their settings.

Setting	Caching mode	Core filer verification	Local Directories	Write Back Delay
Clients Bypassing the Cluster	Read	30 seconds	Disabled	None
Read Caching	Read	Never	Disabled	None
Read and Write Caching	Read/Write	Never	Disabled	1 hour
Full Caching	Read/Write	Never	Enabled	10 minutes
Transitioning Clients Before or After a Migration	Read	30 seconds	Enabled	None

For cloud filers, only Full Caching or a custom policy are enabled.

### Custom Cache Policies

Users can create a new cache policy or modify system default policies from the new **Manage Cache Policies** page. This page replaces the **Cache Policy** core filer pages in previous versions of the Avere Control Panel.

### Expanded Add New Filer Wizard

The **Add New Filer** utility now allows users to set any of the five default caching policies. Choosing a cache policy in the **Add New Filer** dialog automatically shows the full description, diagram, and warnings.

### Changing Cache Policies

Cache policies can be changed from the **Core Filer Details** page. The policy change panel automatically describes the implications and any required actions to change to the new selected policy - for example, changing a policy could cause cluster unavailability while data is moved, and all clients that access the cluster might need to be remounted.

## Suspend a Core Filer or VServer Junction

The new core filer suspend feature allows administrators to disable client access to particular junctions or core filers that are offline or unreachable. This feature can be used to disable a core filer for maintenance, or remove a failed core filer from the system for repair or replacement.

Before this change, Avere OS supported a subset of NFS operations for offline volumes, but this approach could affect performance for other client requests that did not involve the affected core filer. (19967)

Access this feature from the **Settings** tab of the Avere Control Panel.

- To suspend a junction, go to **VServer > Manage VServers**. Select the junction in the **Actions** column and click the **Suspend** button.
- To suspend a core filer, go to **Core Filer > Manage Core Filers**. Select the filer in the **Actions** column and click the **Suspend** button.

## FlashMove Status Indicator

A status tracker now appears in the Avere Control Panel to show the progress of a FlashMove operation. (20652)

## Directory Services Configuration Wizard

A new configuration utility walks users through setting up directory services for a new FXT cluster. The **Directory Services** page (in the **Cluster** menu) has been updated to reflect these changes and to incorporate the settings that previously appeared on the **Active Directory** settings page. (19979)

## vFXT Failover in Multiple AWS Availability Zones

Avere OS now supports deploying vFXT servers in more than one Amazon Web Services Availability Zone (AZ). This feature provides redundancy: if one AZ becomes unavailable, clients can connect to a vFXT in a different AZ.

One vFXT is supported per zone, and up to three zones can be used.

Each vFXT's AWS IAM profile should provide the following permissions:

```
"ec2:AssignPrivateIpAddresses",  
"ec2:UnassignPrivateIpAddresses",  
"ec2:DescribeInstance*",  
"ec2:DescribeRouteTables",  
"ec2:ReplaceRoute",  
"ec2:CreateRoute",  
"ec2>DeleteRoute"
```

This configuration allows the operating vFXTs to reprogram routes and allow clients to connect to the same IP address used on the failed vFXT.

## Speed Improvement when Listing Junctions

The VServer junction listing procedure was enhanced to be more efficient. Before this change, the **Namespace** page in the Avere Control Panel could fail to load for systems with thousands of junctions. (20189)

## Support for Thales eSecurity KMIP Server

This release includes enhancements that allow the Avere OS KMIP server to read encryption keys from Thales eSecurity's keyAuthority server. (20320)

## Expanded Support for Reading ACL Shares over CIFS

This release includes changes that allow the Avere cluster to use ACLs stored on two additional types of core filers:

- NetApp Data ONTAP 8 in Cluster Mode (07076)
- Hitachi HNAS systems (Consult Avere Global Services for details) (20622)

## Microsoft "Previous Versions" Support

This release adds the ability to browse and open previous versions of a file through Windows Explorer. To show previous versions, right-click on a file name, choose Properties, and click the Previous Versions tab. Note that this feature uses the Avere snapshot system, which is not enabled on NAS core filers unless the "snapshot name" parameter is set. (18803)

## Snapshot Cleaner Improved to Eliminate Slowdown

The procedures for estimating available space and cleaning up snapshots was reworked to eliminate increased latency in client operations during these operations. This change includes adjustments to prevent slowdowns on Cleversafe cloud buckets when polling available space. (20728)

## FlashMove and FlashMirror Efficiency Improvements

Several changes were made to streamline operations and improve the robustness of the FlashMove and FlashMirror features. Specific updates are listed in Resolved Issues, below.

# 2. Resolved Issues

## Cloud Object Store

- 00653 Allow administrators to upload signed SSL certificates in place of self-signed Avere certificates
- 19726 Fixed a bug that caused a crash when aborting a cluster snapshot
- 20143 Improved cache management for read-ahead operations when using cloud storage
- 20198 The code for creating and using an Amazon Web Services cloud core filer was updated after an Amazon change caused problems creating core filers using the AWS US-standard region
- 20230 Fixed an issue that caused the df command to return incorrect values for cloud filers
- 20251 Locking and timeout behavior was changed to solve a problem that could cause newly created cloud filers to fail if they could not reach a license server
- 20261 Fixed a bug in the object cache cleaner that could cause a restart
- 20289 Changed the rate limits for writing to cloud storage to prevent out-of-memory errors that caused restarts of the filesystem and kernel. This issue was most often encountered when the cloud backend was in write-through mode while taking snapshots.
- 20299 Fixed a memory leak that could affect clusters using cloud storage
- 20333 Added a minimum TTL setting to account for low TTL values assigned by Windows DNS servers
- 20443 Fixed an error that caused a restart when trying to remove a particular kind of invalid security certificate
- 20540 Fixed a bug that caused a misleading alert about a stuck directory flush operation
- 20559 Added validation to check for errors on put requests
- 20656 A bug was fixed that caused the wrong return code to be sent to an Amazon Web Services cloud core filer
- 20736 Improved write timing to prevent client latency when creating snapshots
- 20741 Fixed an issue where writing files larger than 16GB to cloud core filers could cause long data flush times
- 20748 Updated the API update with ability to specify a new core filer using Swiftstack and "Existing Avere Data" settings
- 20847 Ensure that snapshots are valid after restarting services
- 20868 Improved efficiency in synchronization calls for cloud filers to eliminate metadata backlog
- 20910 Fixed a problem that prevented the .snapshot directory from being seen by an ls -la command
- 21043 Eliminated an "SSL error: certificate verify failed" error caused by an intermediate CA used in Amazon Web Services' change to SHA-256. OCSP and CRL revocation checks now are supported for both AWS and Google Cloud (GCS and GCE) endpoints.
- 21073 Fixed a bug that caused a long delay in accessing a .snapshot directory

## Filesystem

- 17794 Fixed an issue where UID and GID parameters were not passed between Windows clients and core filers
- 19260 Changed buffer design to eliminate a possible core from a memory allocation bug
- 19471 Fixed a locking issue in the Cache Filesystem (CFS) that caused a filesystem restart
- 19913 Fixed a bug that could cause core filers with hierarchical exports to stop responding after an upgrade or change to volume identifiers
- 19999 Fixed a return code issue that could cause a core filer with a read-only cache policy to stop responding
- 20287 Fixed a crash caused by trying to release the same lock twice (This problem only occurred when always-forwarding was enabled and an HA Barrier in place.)
- 20294 Updated output buffer size calculations to handle large files that are resistant to compression; fixes an IO commit error that sometimes occurred when working with large image files
- 20407 Added the ability to customize migration settings individually as well as globally
- 20411 Increased the maximum number of files flushed at one time to 96 from 64
- 20416 Fixed a problem that could cause a node to stop responding after simultaneously removing storage containers and pulling data from core filers
- 20559 Added validation to check for errors on put requests
- 20577 Improved efficiency for cache cleaning
- 20687 Improved operation management process to prevent hanging operations, which could block junctions
- 20839 Fixed an issue where a stuck setattr could cause a core loop
- 20851 A change was made to prevent a stuck operation when a core filer is configured to forward requests to another node
- 20856 Fixed an issue where truncating deleted blocks would cause a filesystem restart
- 20931 Fixed a rare condition that could cause a node restart (This bug existed only in a limited-access version of Avere OS, 4.6.1.4.C2.)
- 21065 Added a timeout and a manual command to stop a directory populate task that is not completing

## FlashMove/FlashMirror

- 19849 Changed the name format of an intermediate directory used in FlashMove and FlashMirror operations to prevent issues with moves in cascaded clusters
- 20091 Key statistics for mirror tasks persist after a node is rebooted. Before this change, move task statistics were persistent but mirror operation stats were purged.
- 20214 Improved file metadata comparison procedures to prevent unneeded copying when mirroring content
- 20449 Migration performance improvement
- 20627 Eliminates the "Protocol transition is not enabled on vserver" error when attempting FlashMove operation between two cloud servers
- 20661 Eliminated a possible race condition that could cause FlashMove migrations to fail
- 20665 Improved the process for protecting volumes during FlashMove operations; in some circumstances, an inefficiency could cause the FlashMove operation to fail or cause an outage
- 20706 Fixed a problem with FlashMove containers that sometimes caused a repopulation error after moving the same cloud source volume to the same cloud target more than once
- 20929 Fixed a FlashMove error message that referenced the wrong file
- 21070 Fixed a data integrity issue with metadata ops after migration tasks

## General

- 16658 Removed unneeded error messages about exceptions from HAVoter: readBackendVote (These messages appeared only in the system log, which is not routinely used for cluster management.)

- 19387 Fixed a problem with buffer memory consumption when using a 9000 MTU jumbo frame network setting
- 19955 Read/write procedures were changed in the statistics server to prevent a server restart under heavy load due to corrupted metadata files
- 19966 Fixed an issue in rpcbind where multiple RPC calls via UDP could cause a restart or crash
- 20270 Removed a warning about write failures ("The Avere cluster is not meeting its data writeback requirements") to prevent confusion
- 20378 References to the cluster startup configuration file (armconfig.xml) were changed and removed from the Avere Control Panel to avoid confusion. (This file does not provide a backup of the full cluster configuration.) There is no longer an option on the Support tab to generate a configuration file, and the initial cluster setup screen now allows you to specify an "Avere setup file" instead of calling it a configuration file.  
  
Contact Avere Global Services if you need to save cluster configuration information or move it from one system to another.
- 20811 Fixed a recently introduced bug that impacted performance of the multi-cluster dashboard in Avere Control Panel
- 20857 Improved the node.getHardwareInfo API to eliminate unnecessary operations that could cause a race condition and error
- 20865 Modified the boot procedure on vFXT nodes to prevent NTP synchronization errors
- 20867 The server that gathers node performance statistics now returns N/A when no stats are available instead of giving an error message
- 20892/20780 Fixed an issue in local directories settings that could prevent the system from adding new core filers
- 20924 Improved responsiveness in the Avere Control Panel by optimizing data rendering code
- 21083 Fix free write space calculations for HA nodes

## NFS

- 15822 Added a flush operation for the duplicate request cache to prevent cores due to bus errors.
- 17584 Enabled specific setting when adding Isilon core filers to prevent setattr from becoming stuck during flush.
- 20671 Fixed an issue where offline core filer volumes and deleted junctions could cause core filer connectivity problems

## SMB/CIFS

- 18930 Prevents users from using invalid characters in core filer share names when creating, modifying, or using shares
- 20613 Fixed an issue where users connecting to an autoloaded home service could cause a core loop
- 20622 Updated code to allow creation of a junction that points to the root of another cluster  
**Note:** This practice is not recommended. Instead, create individual subdirectories for cascaded junctions to allow more diverse core filer types.
- 20742 Added a timeout to fix an uncommon case where a client disconnect could cause a read operation to loop
- 20785 Ensured that only enabled user accounts are returned from Active Directory systems, regardless of what method is used to disable user accounts
- 20883 Fixed a system cache overflow problem related to a bug in secure ID lookup

## Security

- 19879 Updated settings to improve security: enabled server request timeout (10 seconds), and disabled obsolete encryption methods RC4 and SSLv3

# New in Avere OS 4.6.1.4

## 1. New Features and Enhancements

### Support for Google Cloud Storage Nearline Buckets

This release adds support for Nearline storage buckets as archive resources for Google Cloud Storage filers. Avere OS now supports all three Google Cloud Storage classes: Standard, Durable Reduced Availability (DRA), and Nearline.

Because accessing Nearline storage can incur data transfer fees, you must explicitly enable Nearline storage in the Avere Control Panel. Check the Use Nearline Bucket box in the Add New Core Filer screen when configuring an archive bucket.

**Add New Core Filer**

**Creating New Cloud Filer**

Service type: Google Cloud Storage (S3)

Cloud credential: loan19-g1-cred (s3)

Bucket contents:  Empty  Existing Avere data

Use nearline bucket:

Please note that a cloud filer using a nearline storage bucket can result in excessive storage charges due to directory metadata updates. Check the documentation for further information on this topic.

Use HTTPS:

Compression mode: LZ4

Snapshot policy: None

Page 3 of 4

Back Next

### SwiftStack Support

Avere OS now supports SwiftStack cloud storage products as core filers.

### Additional Google High Memory VM Type Supported

Google Compute Engine high memory n1-highmem-32 VMs now can be used as cluster nodes. Avere OS also supports n1-highmem-8 instance types.

## 2. Resolved Issues

### Filesystem

- 18664 Enable optimization to prevent latency when simultaneously writing and listing directory contents
- 19991 Fixed a cache settings issue that could cause FXT restarts
- 20255 Fixed a REST translation problem that caused the FXT to restart

- 20266 Fixed an issue with FlashMove/FlashMirror out-of-sync queue
- 20434 Bug fix to update the cloudEndpoints.xml file after creating or modifying cloud regions
- 20540 Resolved a problem that caused a “flush of directory op is stuck” alert when creating directories
- 20582 Fixed a network lock manager issue that caused stalled operation requests on read-only caches
- 20605 Resolved a lock priority issue that could cause deadlocks

## General

- 19759 Use fast reformat on GCE-connected SSDs to prevent cache errors after node reboot or failure
- 20416 Fixed a deadlock problem that could cause local directory manager problems when simultaneously removing containers and bringing up cold store
- 20197 Enabled access logging per client core for accounting and billing purposes

## NFS

- 20147 Introduced ignoreRootPathExport setting to avoid conflict with NetApp systems running ONTAP 8.2.3

## CIFS

- 1301 Eliminated duplicate connection listings during long writes
- 17794 Fixed an issue where UID and GID parameters were not passed between Windows clients and core filers

## GUI

- 20549 Confirmation dialog appears consistently to confirm mismatched export policies for FlashMove/FlashMirror

# New in Avere OS 4.6.1.3

## 1. New Features and Enhancements

### Enhancements

- 19425 Implemented a faster way to process symlinks for CIFS
- 19655 If the cluster detects mtime change, fetch the file via an async special read ahead thread
- 20091 The FXT will now capture and retain FlashMirror statistics while the core filers are synchronized
- 20185 Accelerated mass.getAll xmlrpc call when querying core filers involved in FlashMove/Mirror operations
- 20270 Removed false positives from “data writeback requirements” alert



## 2. Resolved Issues

### Filesystem

- 20177 Improved performance by breaking parent/child dependency chains under safe conditions
- 20318 Fixed issue where upgrade from 4.5.0.4 C11 to 4.5.1.2 would cause filesystem restarts and core dumps
- 20179 Reduced cloud latencies by minimizing the amount of PUT commands when objects have been cleaned
- 20198 Updated AWS endpoints to match AWS changes to their us-east-1 (us-standard) region
- 20251 Fixed an issue where adding a cloud core filer might result in a filesystem restart.
- 20333 Fixed issue where DNS TTL of 0 would cause address resolution failure.
- 19892 Fixed issue where stopping a mirror and removing a target file would cause a stuck op
- 20207 Fixed an issue where a FlashMove/Mirror job would not complete due to repopulating objects in the out-of-sync queue
- 20214 Fixed a FlashMove/Mirror issue where files would be copied even when they existed in the destination.
- 20002 Fixed issue where deleted file handle pointers had associated inodes and recycling the fhp caused an error loop
- 20312 Reincorporated auxiliary group IDs for cold reads.

### General

- 19726 Fixed an issue where aborting a cloud snapshot would cause a segmentation fault
- 20323 Added a check to prevent starvation-related core loops

### Security

- 19879 Patched security issues on RC4 cipher, HTTPD, and SSLv3 components

### NFS

- 17584 Enabled allowCifsSetattrMode by default for CIFS-ACL NetApp and Isilon core filers with Local Dirs

### CIFS

- 11579 Added share generation number to detect share security style changes and to force client reconnect
- 18930 Share creation now validates name for invalid characters like “/” and reports error
- 19675 Removed multiple component lookups for SMB dirents in order to perform dirents more efficiently
- 19786 Fixed issue where winbindd could cause core dumps and memory leaks
- 20199 Allow non-ASCII characters in domain names



# New in Avere OS 4.6.1.2

## 1. New Features and Enhancements

### Major Features

[AOS-253] The FXT or vFXT cluster now accepts updated token credentials via a new XML-RPC API call. Current support for AWS-style credentials as presented in their Security Token Service (STS).

### Other Enhancements

- 5435 Allow additional outstanding operations running for a single NFS client
- 17711 User/group download from trusted domains
- 19397 Easily edit confdb xml file for the use of dbrestore.sh
- 19425 Added fhp locking that could run by various threads; vcm-cifs readdirplus fast path symlink processing
- 19449 Permit native identity to specific core filers for all users without regard to uid/gid mappings
- 19613 Vcm readdirplus returns last cookie in arg.rsp.lastCookie to assist cifs
- 19655 Proactive-read support via async read-ahead module
- 19952 Support for Commercial Cloud Services (C2S)

## 2. Resolved Issues

### Filesystem

- 17293 Problem finding zero block in vcmmove while running multiops & subdir migrations
- 17372 Solution to stuck migration with Isilon-like exports that don't have dot-dot ('..') pointing to exports themselves
- 19330 Avoid hitting FlushAttrs Race for new objects; not setting fhp's flushAttrsCounter in the newfile change
- 19350 CDirChange deal with inodes not initialized by dirmgr
- 19352 To prevent armada\_main restart, changed CFS free list to mean that an allocation can complete, even if a compress is required, rather than allowing a page to appear in that free list only if an allocation can complete without requiring a compress of the inode page's file handle space
- 19396 Vcm-fhnamelist: make asserts vcm\_osp\_asserts to prevent panic
- 19457 Aborting a migration to cloud may cause filehandle collision; make sure that cloud target of any migration is a wrapped container so abort won't cause filehandle collision
- 19496 Handle the migration for correct massid properly in fsinfo call response for cloud
- 19499 Perpetual readdir jukebox loop on mass with read-only cache mode, check attrs, local dirs and many files
- 19636 Handle possible invalid entry in a dirmgr router cache to prevent node restarts
- 19704 Removed overloaded \_envp field which caused stuck move jobs with Google bucket as a source
- 19733 Fix to cloud migration so it always moves existing objects to pen and verifies an existing target object
- 19743 Validate wrong-index-block reporting, use cdir\_debug dbutils instead of vcm, cleanup

- 19745 Handle invalid assert when path truncated
- 19747 Now get fdir lock in a case where we have child change but no parent dir change to prevent dirmgr panic
- 19751 Invalidate readdir when file handles are invalid
- 19776 Do special cold Read-ahead for first read at off=0 to some random off number to avoid filesystem restarts
- 19783 Ensure that barrier check is done in sync patch in vcm common op init. And check for cloud fhtag in barrierup function call, so that we block cloud ops when the barrier is up.; VCM operation start is incorrectly processing barrier state and causing long-running HA barrier events
- 19785 Increase the clfs version DB age threshold from 2 to 4 hrs to avoid stale file handles on Google core filers
- 19802 Fix container rename process to handle an Isilon root export case

## General

- 13723 Add stalefile for kwatchdog; Update kwatchdog to include less false positives for mbuf shortages
- 19275 Do not run space tracker or snap cleaner if no snapshot has ever been taken.
- 19384 Add snd packets, snd bytes, and rexmit bytes to TCP\_INFO for ActiveSockets/PassiveSockets datadump
- 19436 Enable callback dedup; Fix error in decrementing auto assign off count back to 0 during dispatcher/allocator startup; Split the global and mp newset lock into multiple locks (one per size bucket); Change mmap calls from allocator to use a single stack which is aligned with superpage sizes
- 19640 Bug fix in clfsSnapshotSchedule and remove problematic asserts
- 19657 Fix possible false alert and add knob to optionally disable the alert for slider; make the conformance query aware of the UNKNOWN condition so it won't falsely send out alert
- 19672 fixed the crash point where we check the filehandle state and if invalid we fail with jukebox instead of crash

## SMB/CIFS

- 15854 Non-wildcarded find must use single component lookup
- 19018 SMB2 parameter optimization and fixed issue with home shares
- 19664 Enabled cluster.disableSpecialCaseCifsUpgradePath knob to disable CIFS upgrade code by default
- 19675 wildcarded smb\* find operations for cifs acl shares must avoid per dirent lookup/getfacl
- 19758 Fixed smbaid memory leak by implementing new version of OpenSSL code
- 19765 Using single component find for non-wildcarded lookup

# New in Avere OS 4.6.1.1

## 1. New Features and Enhancements

### Major Features

- [AOS-124] - Expose Disk Utilization Statistics through Avere Control Panel
- [AOS-135] - Add cache space capacity and usage information to Avere Control Panel (Dashboard)
- [AOS-173] - vFXT for Google Compute Engine
- [AOS-154] - Core Filer support for Google Object Storage

## Other Enhancements

- 9090 Need to expose disk utilization stats
- 16479 add two new dashboard point-in-time tables that show file counts and space
- 16651 remove master key version number requirement from the GUI
- 18133 vFXT should support GCE
- 18378 GUI changes to enable FIPS mode
- 18493 VCM's async-write feature is enabled only on AMI systems; off for rest
- 18526 Make stats histogram output a little more usable
- 18527 Readdirplus/Rcc-cache perf: On readdirplus finding inval dirent attrs avoid overwriting dir contents if no mtime change
- 18705 When configuring a new junction against a single-export core filer - default to selecting "/" as the export
- 18726 expose the forceNodeDiscovery command on the GUI "FXT Nodes" page and in the API
- 18835 Enhanced vFXT shepherd
- 19029 add a uniquifier to EC2 IAM role names
- 19102 vcm-perf: proactive-getattrs for negative lookups
- 19104 corefiler workflow must be updated to support the use of KMIP server configurations
- 19109 vcm-readdirplus-perf: cold readdirplus might not attach fhp to ip
- 19123 vcm-read: special cloud marker for throttling handling
- 19126 read-ra-cloud-perf: stop rolling-ra threads after much time passed since last client read
- 19139 Avoid accidental vserver removes
- 19203 Add a note field to the kmip key register request.
- 19215 Add more options to persistent shepherd "create cluster" menu
- 19217 wording of an alert should be changed from "can no longer access" to "cannot access"
- 19296 support gs: URL support when running in GCE
- 19330 vcm-cold-attrs-perf: avoid hitting so-called FlushAttrs Race for new objects
- 19331 vcm-cifs-readdirplus-special-Thread: don't do access-call
- 19337 support "add node" functionality in persistent shepherd
- 19340 GCE vFXT should use SSD for cache
- 19361 Ability to flush whole access cache
- 19380 rrd backoff when the same error is repeating

19405	All flavors of cloud should be covered by one license type
19419	gui cifs local group support
19425	vcm-cifs readdirplus fast path symlink processing
19485	Add disk IO stats table to node details page
19576	Enable vcm-unstable writes for Google vFXT

## 2. Resolved Issues

### Filesystem

19068	smbd memory leak
19071	if kptcd fails due to misconfigured pt/cd then treat as terminal error rather than retrieable
19086	word temporary file not deleted after word exits
19097	Process start times change if kern.boottime changes (e.g. via settimesofday())
19110	Resolve decoupling of samba LastDir and nfs3c cwd
19577	cifs acl share open must check for delete child permission in parent acl if child acl does not contain delete permission
15158	After wiping/upgrading to recover from another issue, filesystem restart (status == VDisk::ERR_DIRMGR_NOT_FOUND...) in ReplayTask::endTask
17025	repopulating cloud core filer after migration leads to fh aliasing
17395	armada_main: PANIC _status == CfsDirMgrLog::ERR_OK    _status == CfsDirMgrLog::ERR_LOGFULL (DirMgrFlushDirLog.cc line=366)
18256	While restarting after hitting armada_main - PANIC entry.second.timeLeft > _starvationCheckInterval
18328	Mknod does not persist type in FDLR correctly due to regression
18400	DirMgr locked on replay records after mass remove
18664	directory listings take extremely long to be returned to clients when there is file modification activity being generated against another node
18677	Failure in ClfsLockSet::regrabLocalLock while writing back files to amplidata
18714	shami failures on _inCommitting panic
18720	Filesystem restart loop in dirmgr SetMipAttrs::applyAttrs
18777	Buffer Overrun that overwrites the magic number of the inode (2)
18941	make DR locks more efficient
19082	crash in DirMgrRouter::MkdirCtx::dirMgrResponse

19099 slow write workloads due to pessimistic reservation in the cloud code

19140 chgrp failure

19164 Long-running truncates - part 2

19176 vcm-read-mode-perf: first read after write after a create may cold-read

19181 Data migration fails with error 6 - exist

19216 removing 2TB file trips "\_entryCount < \_maxObs" assert

19232 removing 2TB file results in 90M LIST calls to amazon

19263 vFXT: nvramnull driver updates

19270 armada\_main core: VcmMultLookupReq::multLookupComponentStatus1 (vcmmultlookup.cc:402)

19272 vcm-readdirplus duplicate cookie issue

19284 Armada Restart (hashIndex != AfhCacheEntry::\_HashIndexInvalid)

19285 vfxt/Google crash due to wantClean flag missing

19326 cifs-casei related crash on create opforward response path

19327 vcm-cifs-multlookup not handling extended-groups support

19333 vcm-handle backend remapping of groups greater than MAXGID (16)

19356 cdir-size-corruption handling: fail instead of asserting

19357 vcm-readdirplus: better handling of tokenmgr's "TOO MANY TOKENS" errs

19368 Segfault in getTreeCallback

19379 Make prune to work with NFS mass in smooth transition mode

19381 vFXT add node scripting need better IP address chooser

19382 While running multiops with large files (7GB), armada\_main bus error in TokenList::ForEachWithOptionalDataRange

19400 VcmReadWorker::startPeekReadWorker

19401 Kernel Panics (Watchdog timeout)

19410 vfxt\_cluster\_util.py script FATAL without letting the user know the reason for the failure

19414 While running multiops against AWS and shoot down nodes, filesystem restart (\_reserveDone == 1)

19481 vcm-readdirplus: cfs buffer overflow on

19484 vcm-readdirplus: dirmgr-write-around issue when no valid attrs

19507 vcm-flush: so called cloud FIRSTWRITE flag was reused by same worker

19510 missing directory entry

- 19512 excessive flushing of parent dir during renames and creates
- 19549 Prevent user from doing WB/WT to WA to WB/WT
- 19554 implement ILE for object modification op which creates FDLR (Read was repeatedly jukeboxing)
- 19561 Readdirplus causes pending operations
- 19563 While running cloud\_verify.py, armada\_main - PANIC \_entryCount < \_maxObs (ofilecache.cc line=2761)
- 19566 Isilon 7.2 missing attributes returned lookup operations ..
- 19654 cdir .snapshot may cause missing dirent

## General

- 16792 Isf reports TCP/UDP connections for each thread
- 17949 /var full led to confdb inconsistencies
- 18404 Panic String: page fault resulting in a kernel core due to MTU confusion on separate management network
- 18922 Fault 100: 'Invalid node name given for scope'
- 19077 mgmtd restart during upgrade from 4.5.0.1\_EA to 4.5.0.2
- 19113 Parsing error with ampersand in Bind DN field - not well formed
- 19168 SPSAgent restart due to bad support setting being saved
- 19195 Export iostat (devstat) performance counters through armada
- 19222 Condition should be raised when drives in a node's aggregate are not all the same size.
- 19226 syslogd stops logging after newsyslog rollover
- 19247 mgmtd panic loop: in DBOwner::ApplyAndDistribute
- 19280 RemoteCommand failure when cluster is configured to use an InternetVLAN on a fib other than 0
- 19343 Partial trace upload time reverts to current time if too much time passes before submit is pressed
- 19366 cluster.proxyConfig not created on upgrade
- 19375 Cloud snapshot schedule issue
- 19393 vFXTs must have larger per-socket buffering
- 19411 Out-of-memory failures
- 19503 Error seen on dashboard
- 19504 gui: incomplete error message in directory services page
- 19506 Make DT\_rowId a unique id without special chars

- 19632 GCE performance improvements
- 19651 armada\_main: (malloc) Error in malloc(): out of memory error
- 19690 Warn that Invalidate WILL destroy data
- 19708 Cloud proxy not taking effect unless armada restarted

## NFS

- 18262 MassVolumeCheck must only gate alerts based on nfs opaque 10
- 19063 Memory leak in NFS
- 19065 add condition when NfsExportTask getattrs to dirmgr fail
- 19130 extended group access check performance
- 19316 nfs3c/smbd auth sys credentials with large group membership must not be limited to 16 by default
- 19325 client specified delete on close must not be forwarded to core filer during native identity creates

## SMB/CIFS

- 19068 smbd memory leak
- 19071 if kptcd fails due to misconfigured pt/cd then treat as terminal error rather than retrievable
- 19086 word temporary file not deleted after word exits
- 19097 Process start times change if kern.boottime changes (e.g. via settimesofday())
- 19110 Resolve decoupling of samba LastDir and nfs3c cwd
- 19165 kptcd with transitive trust
- 19577 cifs acl share open must check for delete child permission in parent acl if child acl does not contain delete permission

# Contact Support - Avere Global Services

Support can be reached by web, phone, or email.

**By web:** <http://www.averesystems.com/support>

**By phone:**

1-888-88-AVERE, Option 2 (Toll-Free)

1-412-894-2570, Option 2

**By email:** [support@averesystems.com](mailto:support@averesystems.com)